Chapter 1

Distributed Optimization

William Macready[1]
David Wolpert[2]
NASA Ames Research Center,
MS 269-4, Moffett Field, CA, 94035

We demonstrate a new framework for analyzing and controlling distributed systems, by solving constrained optimization problems with an algorithm based on that framework. The framework is an information-theoretic extension of conventional full-rationality game theory to allow bounded rational agents. The associated optimization algorithm is a game in which agents control the variables of the optimization problem. They do this by jointly minimizing a Lagrangian of (the probability distribution of) their joint state. The updating of the Lagrange parameters in that Lagrangian is a form of automated annealing, one that focuses the multi-agent system on the optimal pure strategy. We present computer experiments for the k-sat constraint satisfaction problem and for unconstrained minimization of NK functions.

1.1 Introduction

Recently a new framework for analyzing and controlling distributed systems has been developed [6, 7, 8]. This framework starts with a parameterized class of probability distributions, $\mathcal{Q}$, across the joint state of the variables of the system. A Lagrangian function of $q \in \mathcal{Q}$, is minimized to determine a q over the variables of the distributed system. We consider the special case of this probability Lagrangian framework in which $\mathcal{Q}$ is the set of product distributions.

A strength of the framework is the connections it makes to relate disciplines to one another. For example, it can be motivated by using information theory to relate bounded rational game theory to statistical physics [6, 7]. In a noncooper-

[1] wgm@email.arc.nasa.gov
[2] dhw@email.arc.nasa.gov

ative game the agents are statistically independent at any stage of the game, with each agent i choosing its move x_i by sampling its probability distribution (mixed strategy) at that instant, $q_i(x_i)$; the distribution of the joint-moves is a product distribution $q(x \in X) = \prod_i q_i(x_i)$. Inter-agent coupling occurs indirectly, across time, via the updating of the $\{q_i\}$ at the end of each stage. Information theory shows that the (bounded rational) equilibrium of the game is the q optimizing an associated Lagrangian $\mathcal{L}(q)$.

For some games the optimal $q \in \mathcal{Q}$ is the minimizer of the Kullback-Leibler (KL) distance to a distribution p, $D(q\|p) \equiv \sum_x q(x) \ln\big(q(x)/p(x)\big)$ [1], where p is one of the variants of the canonical ensemble of statistical physics. In other words, the Lagrangian in such cases is $D(q\|p)$ for an associated p from statistical physics. In particular, for $\mathcal{Q}$ being the set of product distributions, the bounded rational equilibrium of the game is a mean-field approximation to p.

When the agents share the same utility function $-G(x)$, the optimizer of $\mathcal{L}(q)$ is the distribution that minimizes the expected value of G, subject to any provided constraints and to an overall entropy value that sets the rationalities of the agents. Moreover, the updating of the q_i at the end of each stage of the game can be designed to be a search process for an optimal q. For example, since q is a vector in a Euclidean space, the search can be done with continuous techniques like gradient descent or Newton's method — even if X is a categorical, finite space. Under such updating, the evolution of the game serves as a distributed constrained optimization algorithm. Note how this contrasts with most stochastic optimization algorithms (e.g. simulated annealing). Those algorithms use probability distributions to help guide search for points $x = [x_1, \cdots, x_N] \in X$ optimizing a function $G(x)$. In contrast, we search over distributions directly.

In many optimization problems, particularly Constraint Satisfaction Problems (CSPs), we want to find multiple solutions q. Multiple runs of the game outlined above might not find different q. Here we show how to construct a single game to obtain multiple distinct solutions at once. The approach is to reparameterize X so that a product distribution over the new parameters corresponds to a coupled distribution across X. We consider such a reparameterization that results in a mixture of M product distributions $q(x) = \sum_m q_0(m) q^m(x)$ [4]. As described below, the associated Lagrangian "pushes" the separate $q^m(x)$ apart.

We begin in §1.2 by elaborating our Lagrangian for mixture models, and consider simple methods to minimize this Lagrangian in §1.3. Experimental validation is presented for k-satisfiability (§1.4.1) and NK (§1.4.2) problems.

1.2 Specifying the Lagrangian

To specify the Lagrangian we must first fix the distribution $p(x)$ we wish to get as close (in KL distance) to. If the objective function we wish to minimize is $G(x)$ (i.e., G is the negative of the utility shared by the bounded rational agents) then we consider the T-parameterized Boltzmann distribution $p(x) = \exp\big(-G(x)/T\big) \;/\; Z(T)$ (At low T — high rationalities — this distribution is concentrated on x having low G values.) The KL distance to this p is

proportional to

$$\mathcal{L}(q) = E_q[G] - TS(q) \tag{1.1}$$

where $E_q[G] \equiv \sum_x q(x)G(x)$, and $S(q) \equiv -\sum_x q(x) \ln q(x)$ is the entropy of q. For product q's $S(q) = \sum_i S(q_i)$ where $S(q_i) = -\sum_{x_i} q_i(x_i) \ln q_i(x_i)$.

Since we are interested in problems with constraints, it is natural to write $G(x) = O(x) + \sum_a \lambda_a c_a(x)$ where O is an objective to be minimized, and the c_a are a set of constraint functions that are required to be non-negative. The λ_a are Lagrange multipliers that are used to enforce the constraints. (In CSP's $O(x) = 0$.)

As mentioned above, we consider distributions of the form $q(x) = \sum_{m=1}^{M} q_0(m)q^m(x)$ where $\sum_m q_0(m) = 1$ and $q^m(x) = \prod_i q_i^m(x_i)$. Substituting this into (1.1) gives the mixture Lagrangian

$$L(q) = \sum_m q_0(m)E_{q^m}[G] - TS(q) = \sum_m q_0(m)\mathcal{L}(q^m) - TJ(q) \tag{1.2}$$

with $\mathcal{L}(q^m)$ given by (1.1) and $J(q) \geq 0$ being the Jensen-Shannon (JS) distance,

$$J(q) = S\Big(\sum_m q_0(m)q^m\Big) - \sum_m q_0(m)S(q^m) = -\sum_m \sum_x q_0(m)q^m(x) \ln \frac{q(x)}{q^m(x)}.$$

The JS term pushes the optimal q^m to differ from each other. Unfortunately, it also couples all variables (because of the sum inside the logarithm), preventing a distributed solution. Thus, we replace J with another function which lower-bounds J and which requires less communication between agents.

A Variational Lagrangian

Following [2], we introduce M variational functions $w(x|m)$ and lower-bound the true JS distance with

$$\begin{aligned} J(q) &= -\sum_m \sum_x q_0(m)q^m(x) \ln \left[\frac{1}{w(x|m)} q_0(m) \frac{w(x|m)q(x)}{q_0(m)q^m(x)}\right] \\ &= \sum_m \sum_x q_0(m)q^m(x) \ln w(x|m)) - \sum_m q_0(m) \ln q_0(m) \\ &\quad - \sum_m \sum_x q_0(m)q^m(x) \ln \frac{w(x|m)q(x)}{q_0(m)q^m(x)}. \end{aligned}$$

Now replace M of the $-\ln$ terms with the lower bound $-\ln x \geq -\nu x + \ln \nu + 1$ obtained from the Legendre dual of the logarithm to find

$$\begin{aligned} J(q) \geq J(q, w, \nu) \equiv &\sum_m \sum_x q_0(m)q^m(x) \ln w(x|m) - \sum_m q_0(m) \ln q_0(m) \\ &- \sum_m \nu_m \sum_x w(x|m)q(x) + \sum_m q_0(m) \ln \nu_m + 1. \end{aligned}$$

We optimize over w and ν to maximize this lower bound. To further aid in distributing the algorithm we restrict the class of variational $w(x|m)$ to products: $w(x|m) = \prod_i w_i(x_i|m)$. For this choice

$$J(q, w, \nu) \equiv \sum_m q_0(m) \left\{ B^{m,m} - \sum_{\tilde{m}} A^{m,\tilde{m}} \nu_{\tilde{m}} + \ln \nu_m \right\} + S[q_0] + 1 \qquad (1.3)$$

where $A_i^{\tilde{m},m} \equiv \sum_{x_i} q_i^{\tilde{m}}(x_i) w_i(x_i|m)$, $A^{\tilde{m},m} \equiv \prod_{i=1}^d A_i^{\tilde{m},m}$, $B_i^{m,m} \equiv \sum_{x_i} q_i^m(x_i) \ln w_i(x_i|m)$, and $B^{m,m} \equiv \sum_{i=1}^d B_i^{m,m}$.[3] At any temperature T the variational Lagrangian which must be minimized with respect to q, w and ν (subject to appropriate positivity and normalization constraints) is then

$$L(q, w, \nu) = \sum_m q_0(m) \mathcal{L}(q^m) - T J(q, w, \nu). \qquad (1.4)$$

1.3 Minimizing the Lagrangian

Even if $x \in X$ is a discrete quantity (as in the cases we consider here) the optimization variables q determined by minimizing L for a fixed λ are continuous so that gradient methods may be applied. Optimizing for the variational parameters w and ν we find

$$\frac{1}{\nu_m} = \frac{1}{q_0(m)} \sum_{\tilde{m}} q_0(\tilde{m}) A^{\tilde{m},m}. \qquad (1.5)$$

$$w_i(x_i|m) \propto \frac{q_0(m) q_i^m(x_i)}{\nu_m} \left[\sum_{\tilde{m}} q_0(\tilde{m}) q_i^{\tilde{m}}(x_i) \frac{A^{\tilde{m},m}}{A_i^{\tilde{m},m}} \right]^{-1}. \qquad (1.6)$$

The dependence of L on $q_0(m)$ is particularly simple: $L(q, w, \nu) = \sum_m q_0(m) \mathcal{E}(m) - T\big(S(q_0) + 1\big)$, where

$$\mathcal{E}(m) = E_{q^m}(H) - T \left(S[q^m] + B^{m,m} - \sum_{\tilde{m}} A^{m,\tilde{m}} \nu_{\tilde{m}} + \ln \nu_m \right),$$

Thus the mixture weights are Boltzmann distributed:

$$q_0(m) = \frac{\exp\big(-\mathcal{E}(m)/T\big)}{\sum_{\tilde{m}} \exp\big(-\mathcal{E}(\tilde{m})/T\big)}. \qquad (1.7)$$

The determination of $q_i^m(x_i)$ is similar. The relevant terms in L involving $q_i^m(x_i)$ are $L \approx q_0(m) \sum_{x_i} \mathcal{E}_m(x_i) q_i^m(x_i) - T S(q_i^m)$ where

$$\mathcal{E}_m(x_i) = E_{q^m_{\setminus i}}(H|x_i) - T \left(\ln w_i(x_i|m) - \sum_{\tilde{m}} \frac{A^{m,\tilde{m}}}{A_i^{m,\tilde{m}}} \nu_{\tilde{m}} w_i(x_i|\tilde{m}) \right).$$

[3] Note that if $w_i(x_i|m) = 1/|X_i|$ is uniform across x_i then $A_i^{\tilde{m},m} = 1/|X_i|$ and $B_i^{m,m} = -\ln |X_i|$. Maximizing over ν_m we find that $J(q, w = 1/|X|, \nu = \nu^*) = 0$. Thus, maximizing with respect to w increases the JS distance from 0.

The conditional expectation $E_{q^m_{\setminus i}}[G|x_i]$ is $\sum_{x_{\setminus i}} G(x_i, x_{\setminus i}) q^m_{\setminus i}(x_{\setminus i})$ where $x_{\setminus i} \equiv [x_1, \cdots, x_{i-1}, x_{i+1}, \cdots x_d]$ and $q^m_{\setminus i}(x_{\setminus i}) = \prod_{j=1(\neq i)}^{d} q_j(x_j)$. The mixture probabilities are thus determined as

$$q^m_i(x_i) = \frac{\exp\bigl(-\mathcal{E}_m(x_i)/T\bigr)}{\sum_{x_i} \exp\bigl(-\mathcal{E}_m(x_i)/T\bigr)}. \tag{1.8}$$

Note that these results requires minimal communication between agents. Assign a 0 agent manage the determination of $q_0(m)$ and (i, m) agents to manage determination of $q^m_i(x_i)$. The M (i, m) agents for a fixed i communicate their $w_i(x_i|m)$ to determine $A^{m,\tilde{m}}_i$. These results along with the $B^{m,m}_i$ from each (i, m) agent are then forwarded to the 0 agent who forms $A^{m,\tilde{m}}$ and $B^{m,m}$ broadcasts this back to all (i, m) agents.

Updating Lagrange Multipliers: In order to satisfy the constraints we must also update the Lagrange multipliers. To minimize communication between agents this is done in the simplest possible way – by taking the partial derivatives with respect to λ. This gives

$$\lambda_j \leftarrow \lambda_j + \alpha_\lambda E_{q^*}[c_j(x)] \tag{1.9}$$

where α_λ is a step size and q^* is the minimizer of $\mathcal{L}$ at the old settings of the multipliers.

Estimation of Conditional Expectations: All update rules require estimation of conditional expectations with some variables clamped to particular values. These are estimated exactly if a closed form expression is available, or with Monte-Carlo sampling if no simple closed form exists. For the problems addressed here the expectations may be evaluated closed form, but Monte Carlo sampling can also be used [6, 8].

1.4 Experiments

We test the method on two different problems: a k-sat constraint satisfaction problem having multiple feasible solutions, and an unconstrained optimization of an NK function.

1.4.1 *k*-sat

The k-sat problem is perhaps the best studied CSP [5]. The goal is to assign N binary variables x_i values so that C clauses are satisfied. The ath clause involves k variables labeled by $v_{a,j} \in [1, N]$ (for $j \in [1, k]$), and k binary values associated with each a and labeled by $\sigma_{a,j}$. The ath clause is satisfied iff $c_a(x) \equiv \bigvee_{j=1}^{k} [x_{v_{a,j}} = \sigma_{a,j}]$ is true. Accordingly we write $G(x, \lambda) \equiv \lambda^\top c(x)$ where λ and c are vectors of length C whose a components are λ_a, and $c_a(x)$ respectively.

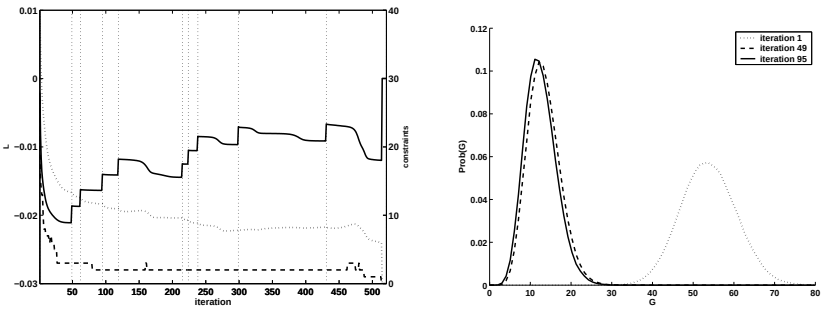

(a) (b)

Figure 1.1: (a) Evolution of Lagrangian value (solid line), expected constraint violation (dotted line), and constraint violations of most likely configuration (dashed line). (b) $P(G)$ after minimizing the Lagrangian for the first 3 multiplier settings. At termination $P(G) = \delta(G)$.

Noting that the ath clause is violated only when all $x_{v_{a,j}} = \overline{\sigma}_{a,j}$ (with $\overline{\sigma} \equiv \text{not}\,\sigma$), the Lagrangian over product distributions can be written as $\mathcal{L}(q) = \lambda^\top c(q) - TS(q)$ where $c(q)$ is the C-vector of expected constraint violations whose ath component is $c_a(q) \equiv \prod_{j=1}^{k} q_{v_{a,j}}(\overline{\sigma}_{a,j})$. The only communication required to evaluate G and its conditional expectations is between agents appearing in the same clause. Typically, this communication network is sparse; for the $N = 100$, $k = 3$, $C = 430$ variable problem each agent interacts with only 6 other agents on average.

For any fixed setting of the Lagrange multipliers, the Lagrangian is minimized by iterating the equations (1.5) – (1.8). Rather than update a single agent at a time we randomly select a subset of variables no two of which impact each other and update the subset simultaneously. The minimization is terminated at a local minimum q^*. If all constraints are satisfied at q^* we return the solution $x^* = \arg\max_x q^*(x)$ otherwise the Lagrange multipliers are updated according to Eq. (1.9). In the present context, this updating rule offers a number of benefits. Firstly, those constraints which are violated most strongly have their penalty increased the most, and consequently, the agents involved in those constraints are most likely to alter their state. Secondly, the Lagrange multipliers contain a history of the constraint violations (since we keep adding to λ) so that when the agents coordinate on their next move they are unlikely to return a previously violated state. Lastly, rescaling the Lagrangian by the norm of λ gives $\mathcal{L}(q) = \hat{\lambda}^\top c(q) - TS(q)/\|\lambda\|$ where $\hat{\lambda} = \lambda/\|\lambda\|$ so that the updating the Lagrange multipliers can be seen as defining a cooling schedule where $T \to T/\|\lambda\|$. The parameter α_λ thus governs the overall rate of cooling. We used $\alpha_\lambda = 0.5$.

Fig. 1.1 presents results for a 100 variable $k = 3$ problem using a single mixture. The problem is satisfiable formula `uf100-01.cnf` from SATLIB

(`www.satlib.org`). It was generated with the ratio of clauses to variables being near the phase transition, and consequently has few solutions. Fig. 1.1(a) shows the variation of the Lagrangian, the expected number of constraint violations, and the number of constraints violated in the most probable state $x_{\rm mp} \equiv \arg\max_x q(x)$ as a function of the number of iterations. The starting state is the maximum entropy configuration, and the starting temperature is $T = 0.0015$. The iterations at which the Lagrange multipliers are updated are indicated by vertical dashed lines which are clearly visible as discontinuities in the Lagrangian values. To show the stochastic underpinnings of the algorithm we plot in Fig. 1.1(b) the probability density of the number of constraint violations obtained as $\text{Prob}(G) = \sum_x q(x)\delta\big(G - G(x,1)\big)$.

Results on a larger problem with more mixtures are shown in Fig. 1.2(a). This is the 250 variable/1065 clause problem `uf250-01.cnf` from SATLIB with the first 50 clauses removed so that the problem has multiple solutions. The initial temperature is 0.1. We plot the number of constraints violated in the most probable state of each mixture as a function of the number of updates. as well as the expected number of violated constraints. After 8000 steps three distinct solutions have been found along with a fourth solution which violates a single constraint.

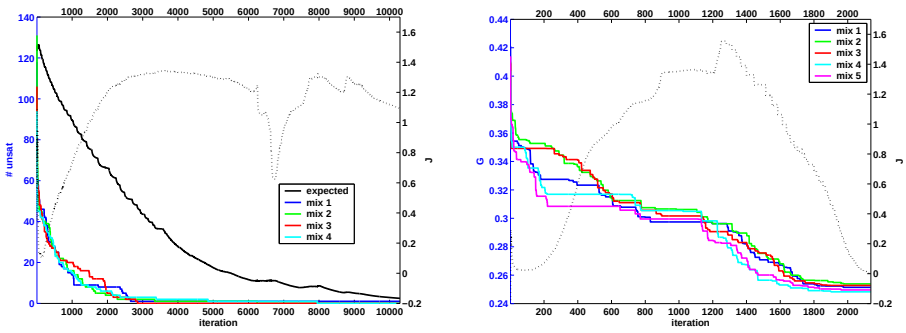

(a) (b)

Figure 1.2: (a) The solid colored curves show the number of unsatisfied clauses in of the best $x_{\rm mp}$ configurations of each of the 4 mixtures vs iterations. The solid black line plots the expected number of violations, and the dashed black line shows the approximation to the JS distance. (b) The solid colored curves show the evolution of the G value of the best $x_{\rm mp}$ configurations for each of 5 mixtures versus number of iterations. The dashed black line shows the corresponding approximation to the JS distance.

1.4.2 Minimization of NK Functions

The NK model defines a family of tunably difficult optimization problems [3]. The energy of N binary variables is defined as the average of N contributions local to each variable x_i and involving K other randomly chosen variables $x_i^1 \cdots x_i^K$: $G(x) = N^{-1} \sum_{i=1}^{N} E_i(x_i; x_i^1, \cdots x_i^K)$. For each of the 2^{K+1} local configurations E_i is assigned a value drawn uniformly from 0 to 1. Fig. 1.2(b) plots the energy of a 5 mixture model for a multi-modal $N = 300$ $K = 2$ function. At termination 5 distinct configurations are obtained with the nearest pair of solutions having Hamming distance 12.

1.5 Conclusion

A distributed constrained optimization framework based on probability Lagrangians has been presented. Motivation for the framework was drawn from an extension of full-rationality game theory to bounded rational agents. An algorithm was developed and demonstrated on two problems. The results show a promising, highly distributed, off-the-shelf approach to constrained optimization.

Bibliography

[1] Cover, T., and J. Thomas, *Elements of Information Theory*, Wiley-Interscience New York (1991).

[2] Jaakkola, T. S., and M. I. Jordan, "Improving the mean field approximation via the use of mixture distributions", *Learning in Graphical Models*, (M. I. Jordan ed.). Kluwer Academic (1998).

[3] Kauffman, S. A., and S. A. Levin, "Towards a general theory of adaptive walks on rugged landscapes", *Journal of Theoretical Biology* **128** (1987), 11–45.

[4] Macready, W. G., and D. H. Wolpert, "Distributed constrained optimization", preprint cond-mat/XXX (June 2004).

[5] Mezard, M., G. Parisi, and R. Zecchina, "Analytic and algorithmic solution of random satisfiability problems", *Science* (June 2002).

[6] Wolpert, D. H., "Factoring a canonical ensemble into a product of ensembles", preprint cond-mat/0307630.

[7] Wolpert, D. H., "Bounded rational games, information theory, and statistical physics", *Complex Engineering Systems* (D. Braha and Y. Bar-Yam eds.), (2004).

[8] Wolpert, D. H., "Generalizing mean field theory for distributed optimization and control", submitted.